

WHAT IS CLAIMED IS:

1. A switching system having an input and an output, the switching system further comprising:

a first communications switch and a second communications switch
5 connected by at least one communications link, comprising at least one channel, for transmitting a plurality of data packets from said communications link to the output of the switching system;

a Common Time Reference (CTR), divided into a plurality of contiguous
periodic super cycles (SCs) each comprised of at least one contiguous time cycle (TC)
10 each comprised of at least one contiguous time frame (TF);

wherein each of the communications switches is further comprised of a
plurality of input ports and a plurality of output ports, each of the input ports connected
to and receiving data packets from the communications link from at least one of the
channels, and each of the output ports connected and transmitting data packets to the
15 communications link over at least one of the channels;

wherein each of the communications links is connected between one of the
output ports on the first communications switch and one of the input ports on the second
communications switch;

wherein each of the communications switches has a switch controller,
20 coupled to the CTR, the respective input ports, and the respective output ports;

wherein each of the communications switches has a switch fabric coupled to the respective switch controller, the respective input ports, and the respective output ports;

5 wherein each of the switch controllers is responsive to the CTR for scheduling connection to the switch fabric from a respective one of the input ports, on a respective one of the input channels during a respective one of the time frames;

10 wherein each of the switch controllers defines the coupling from each one of the respective input ports for data packets received during any one of the time frames, on a respective one of the channels, for output during a predefined time frame to at least one selected one of the respective output ports on at least one selected respective one of the channels; and

15 wherein the data packets that are output during a first predefined time frame on a selected respective one of the channels from the respective output port on the first communications switch are forwarded from the respective output port of the second communications switch during a second predefined time frame on a selected respective one of the channels responsive to the CTR.

2. The system as in claim 1,

20 wherein the plurality of input ports each receives data packets over at least one of a plurality of incoming channels (j), and wherein the plurality of output ports each sends data packets over at least one of a plurality of outgoing channels (l);

wherein each of the incoming channels (*j*) has a unique time reference
(UTR-*j*) that is independent of the CTR;

wherein the (UTR-*j*) is divided into super cycles, time cycles, and time
frames of the same durations as the super cycles, time cycles, and time frames of the
5 CTR;

wherein each of the super cycles, time cycles, and time frames of the
(UTR-*j*) have a start and end in time that is different than the respective start and end in
time of the super cycles, time cycles, and time frames of the CTR.

10 3. The system as in claim 2, further comprising:

a plurality of buffer queues, wherein each of the respective buffer queues
is associated, for each of the time frames, with a unique combination of one of the
incoming channels and one of the outgoing channels; and

a mapping controller within the switch controller system for logically
15 mapping, for each of the (UTR-*j*) time frames, selected incoming channels (*j*) to selected
buffer queues, and for logically mapping, for each of the CTR time frames, selected ones
of the plurality of buffer queues to selected outgoing channels (*l*);

wherein each of the buffer queues is further comprised of an alignment
subsystem comprised of a plurality of time frame queues, wherein each of the time frame
20 queues comprises means to determine that the respective time frame queue is empty,

wherein each of the time frame queues further comprises means to determine that the respective time frame queue is not empty;

wherein the data packets that arrive via the incoming channel (j) are stored in the respective time frame queue of the alignment subsystem responsive to the mapping controller; and

wherein the mapping controller further provides for coupling of selected ones of the time frame queues to respective ones of the outgoing channels (l), for transfer of the respective stored data packets during the respective associated CTR time frames.

4. The system as in claim 3,
wherein the alignment subsystem, responsive to the mapping controller, transfers all of the data packets associated with a respective first time frame as defined by the UTR- j into an empty first time frame queue from incoming channel (j), during the respective selected first time frame of the time frames (TFs) as defined by UTR- j ,
wherein the respective time frame queue is designated as full;

wherein the alignment subsystem, responsive to the mapping controller, transfers data packets out of a full second time frame queue to outgoing channel (l), during a selected one of the time frames (TFs) as defined by UTC, wherein the second time frame queue is designated as empty; and

wherein the first time frame queue and the second time frame queue are mutually exclusive at all times.

5. The switch controller system as in claim 4, wherein the time frame queues are comprised of at least two, three, and more than three time frame queues.

6. The system as in claim 2, wherein the communications link is an optical link with a plurality of optical channels, the system further comprising:
means for adding a delay element to a selected one of the input ports.

7. The system as in claim 6, further comprising:
wherein the delay element provides for phase aligning the UTR-j with the CTR by adding a link delay equal to the difference between a beginning of the respective CTR time frame and a beginning of the respective UTR-j time frame.

8. The system as in claim 6, wherein the delay element provides phase alignment of a start of a respective one of the CTR time cycles relative to a start of a respective one of the UTR-j time cycles.

9. The system as in claim 6, wherein the delay element provides phase alignment of a defined point in a respective one of the CTR time cycles to a defined point in a respective one of the UTR-j time cycles.

10. The system as in claim 6, wherein the delay element is further comprised of a passive optical fiber.

11. The system as in claim 6, wherein the delay element is further comprised of an optical fiber having programmable tap points.

12. The system as in claim 11, wherein the programmable tap points are further comprised of optical switches.

13. The system as in claim 6, wherein each of the input ports is further comprised of an optical receiver, wherein the delay element is a part of the optical receiver.

14. The system as in claim 1, further comprising a switching fabric for coupling the switching system input to the switching system output.

15. The system as in claim 14, wherein the switching fabric is at least one of the following: a crossbar, a generalized multi-stage cube network, a Clos network, a Benes network, an Omega network, a Delta network, a multi-stage shuffle exchange network, a Banyan network, a combination of demultiplexers and multiplexers, and an optical switch.

16. The system as in claim 1,
wherein there are a plurality of the first communication switches;
wherein there are a plurality of the communications links;
where each of the communications links has a plurality of channels, each
5 associated with a respective wavelength.

17. The system as in claim 16, further comprising:
means for coupling a first predefined subset of the channels for each
respective one of the communications links from the respective communications link to a
10 second defined one of the communications links.

18. The system as in claim 17, wherein the respective communications link is
the same as the second defined one of the communications links.

19. The system as in claim 17, wherein the means for coupling is an optical
switch.

20. The system as in claim 19, wherein the optical switch demultiplexes the
first predefined subset into a predefined respective second predefined subset of the
20 respective channels.

21. An input interface system for mapping an asynchronous stream of data packets, each comprising a header portion and a payload portion, from at least one source to at least one destination, said system comprising:

a Common Time Reference (CTR), divided into a plurality of contiguous periodic super cycles each comprised of at least one contiguous time cycles each comprised of at least one contiguous time frame (TF);

at least one synchronous virtual pipe (SVP) having a subset of predefined time frames uniquely associated therewith;

at least one queue wherein each queue is associated with a respective one of the SVPs;

means for analyzing the header portions of the asynchronous data packets; means for storing the analyzed data packets in respective queues responsive to the means for analyzing;

a link coupled to the destination; and

an SVP Forwarding Controller, comprising a second memory for storing SVP schedules, and for forwarding, to the link, respective ones of the asynchronous data packets from respective ones of the queues responsive to the respective SVP schedule and the CTR.

22. The system as in claim 21, wherein there are a plurality of SVPs, and wherein there are a plurality of respective associated queues.

24. The system as in claim 23, wherein the cyclically recurring order is a predefined number of at least one time cycle.

10

15

27. The system as in claim 26, wherein the recurring order starts at an arbitrary point of time in the CTR.

20 wherein the SVP Forwarding Controller provides mapping for forwarding
of the respective data packets from a respective one of the queues to a respective one of

the channels during selected respective ones of the time frames, responsive to the SVP
schedules and the CTR.

29. The system as in claim 28, wherein the SVP Forwarding Controller is
5 comprised of a plurality of SVP forwarding controllers.

30. The system as in claim 29, wherein each of the plurality of SVP
Forwarding Controllers is associated with at least one of the channels.

31. The system as in claim 29, wherein there are a plurality of sets of queues,
10 each set comprising at least one queue, wherein each set is associated with one respective
one of the SVP Forwarding Controllers.

32. The system as in claim 21, wherein there are a plurality of separate and
15 independent streams of asynchronous data packets.

33. The system as in claim 32, wherein there are a plurality of SVP
Forwarding Controllers each associated with at least one of the plurality of asynchronous
data streams.

20

34. The system as in claim 33, wherein there are a plurality of sets of queues, each set comprising at least one queue, wherein each set is associated with one respective one of the SVP Forwarding Controllers.

5 35. The system as in claim 32, wherein there are a plurality of the means for analyzing;

wherein each of the means for analyzing provides analysis of at least one of the plurality of streams of asynchronous data packets.

10 36. The system as in claim 32, wherein there are a plurality of means for analyzing; wherein each of the plurality of streams is associated with at least one of the means for analyzing.

15 37. An input interface system for mapping an asynchronous stream of data packets, each comprising a header portion and a payload portion, from at least one source to at least one destination, said system comprising:

a Common Time Reference (CTR), divided into a plurality of contiguous periodic super cycles each comprised of at least one contiguous time cycles each comprised of at least one contiguous time frame (TF);

20 at least one synchronous virtual pipe (SVP) having a subset of predefined ones of the time frames uniquely associated therewith;

means for analyzing the header portions of the asynchronous data packets;

5 means for storing the analyzed data packets in respective queues

responsive to the means for analyzing;

38. The system as in claim 37,
wherein each of the queues is subdivided into a Constant Bit Rate (CBR),
a Variable Bit Rate (VBR), and a best efforts (BE) queue;

wherein the means for storing provides for storage of the respective data packets in the respective CBR, VBR, and BE queues for an associated respective one of the queues associated with an associated respective one of the time frames, responsive to the means for analyzing.

39. The system as in claim 38,

wherein the output from the respective ones of the queues is prioritized to provide first for output from the respective one of the queue's CBR queue, then from the respective one of the queue's VBR queue, and then from the respective one of the queue's
5 BE queue.

40. The system as in claim 39, wherein in one case, certain ones of the data packets from the CBR, VBR, and BE queues for the respective one of the time frames are not output during a respective associated one of the time frames, the system further
10 comprising:

a rescheduling controller for detecting the one case and for rescheduling the certain ones of the data packets.

41. The system as in claim 40, wherein the rescheduling is provided
15 responsive to the controller in the means for analyzing.

42. The system as in claim 38, wherein in one case during at least one given time frame, certain ones of the data packets from the respective SVP are not output, the system further comprising:

a rescheduling controller for detecting said one case, and providing for rescheduling of the certain ones of the data packets for a next available one of the subset of time frames associated with the respective SVP.

5 43. The system as in claim 37, wherein the data packets are forwarded out of the respective queues during predefined one of the time frames in a cyclically recurring order.

10 44. The system as in claim 43, wherein the cyclically recurring order is a predefined number of at least one time cycle.

 45. The system as in claim 43, wherein the cyclically recurring order is a predefined number of at least one super cycle.

15 46. The system as in claim 43, wherein the cyclically recurring order is a summation of a predefined number of time frames plus a predefined number of time cycles plus a predefined number of super cycles.

20 47. The system as in claim 37, wherein the Forwarding Controller is comprised of a plurality of forwarding controllers, wherein each of the plurality of forwarding controllers is associated with at least one of the channels.

48. The system as in claim 47, wherein there are a plurality of sets of queues, each set comprising at least one queue, wherein each set is associated with a respective one of the forwarding controllers.

5 49. The system as in claim 47, wherein there are a plurality of separate and independent streams of asynchronous data packets.

10 50. The system as in claim 49, wherein there are a plurality of means for analyzing each associated with at least one of the plurality of asynchronous data streams.

11 51. The system as in claim 50, wherein there are a plurality of sets of queues, each set comprising at least one queue, wherein each set is associated with a respective one of the forwarding controllers.

15 52. The system as in claim 51, wherein each of the means for analyzing provides analysis of at least one of the plurality of streams of asynchronous data packets.

20 53. The system as in claim 37,
wherein each of the time frames is associated with at least one data packet;
and

wherein the at least one data packet is at least one of an Internet Protocol (IP) data packet, a plurality of IP data packets, an Multi-Protocol Label Switching (MPLS) data packet, a plurality of MPLS data packets, a frame relay (FR) data packet, a plurality of FR data packets, a Fiber Channel (FC) data packet, a plurality of FC data packets, an Asynchronous Transfer Mode (ATM) cell, and a plurality of ATM cells.

54. The control system as in claim 37,
wherein the CTR is Coordinated Universal Time (UTC) standard; and
wherein the super cycle is one of a single UTC second, a predefined integer number of UTC seconds, and a fraction of one UTC second.

55. A communications system, comprising:
means for mapping an asynchronous stream of data packets, each comprising a header portion and a payload portion, through an input interface system via a communications link from at least one source to at least one switching subsystem;

a Common Time Reference (CTR), divided into a plurality of contiguous periodic super cycles each comprised of at least one contiguous time cycles each comprised of at least one contiguous time frame (TF);

wherein the input interface subsystem is comprised of:
at least one synchronous virtual pipe (SVP) having a subset of predefined time frames uniquely associated therewith;

60. The system as in claim 57, wherein the cyclically recurring order is a summation of a predefined number of time frames plus a predefined number of time cycles plus a predefined number of super cycles.

5 61. The system as in claim 60, wherein the recurring order starts at an arbitrary point of time in the CTR.

62. The system as in claim 55, wherein the link is comprised of at least one of a plurality of channels;

10 wherein the SVP Forwarding Controller provides mapping for forwarding of the respective data packets from a respective one of the queues to a respective one of the channels during selected respective ones of the time frames, responsive to the SVP schedules and the CTR.

15 63. The system as in claim 55, wherein there are a plurality of separate and independent streams of asynchronous data packets.

64. The system as in claim 63, wherein there are a plurality of SVP Forwarding Controllers each associated with at least one of the plurality of asynchronous
20 data streams.

65. The system as in claim 55, wherein the switching subsystem is comprised of at least one input port, at least one output port, and a switching fabric for coupling a respective selected one of the input ports to a respective selected one of the output ports.

5 66. The system as in claim 65, wherein the at least one input port is further comprised of an alignment subsystem for aligning timing for the data packets received at the input port, relative to the CTR.

10 67. The system as in claim 66, wherein the alignment subsystem provides for aligning the timing for each time frame for a set comprising the data packets received at the input port during the respective one of the time frames, the system further comprising:

means for scheduling the transfer of the respective set of the data packets during a subsequent time frame associated with the respective input port.

15 68. A method for switching a plurality of data packets from an input to an output, via a switching system comprising at least a first communications switch and a second communications switch connected by at least one communications link comprising at least one channel, wherein each of the communications switches is further comprised of
20 a plurality of input ports each connected and receiving data packets from the communications link from at least one of the channels, and a plurality of output ports

each connected and transmitting data packets to the communications link over at least one of the channels, wherein each of the communications switches has a switch controller coupled to the input ports and the output ports, wherein each of the communications switches has a switch fabric coupled to the switch controller, the input ports, and the output ports, the method further comprising:

transmitting a plurality of data packets from the link to the output of said switching system;

providing a Common Time Reference (CTR), divided into a plurality of contiguous periodic super cycles each comprised of at least one contiguous time cycles each comprised of at least one contiguous time frame (TF);

coupling the CTR to the switch controller, wherein the switch controller is in part responsive to the CTR;

connecting each of the communications links between one of the output ports on the first communications switch and one of the input ports on the second communications switch;

scheduling connection to the switch fabric from a respective one of the input ports, on a respective one of the input channels during a respective one of the time frames responsive to the CTR;

coupling from each one of the input ports for data packets received during any one of the time frames, on a respective one of the channels, for output during a

predefined time frame to at least one selected one of the output ports on at least one selected one of the channels, responsive to the switch controller; and

forwarding from the output port of the second communications switch during a second predefined time frame on a selected one of the channels, the respective data packets that are output during a first predefined time frame on a selected one of the channels from the output port on the first communications switch responsive to the switch controller.

69. The method as in claim 68,

wherein the plurality of input ports each receives data packets over at least one of a plurality of incoming channels (j), and wherein the plurality of output ports each sends data packets over at least one of a plurality of outgoing channels (l);

wherein each of the incoming channels (j) has a unique time reference (UTR- j) that is independent of the CTR,

wherein the (UTR- j) is divided into super cycles (SCs), time cycles (TCs), and time frames (TFs) of the same durations as the super cycles (SCs), time cycles (TCs), and time frames (TFs) of the CTR;

wherein each of the super cycles (SCs), time cycles (TCs), and time frames (TFs) of the (UTR- j) start and end in time that is different than the respective start and end in time of the super cycles (SCs), time cycles (TCs), and time frames (TFs) of the CTR.

70. The method as in claim 69, further comprising:

providing a plurality of buffer queues, wherein each of the respective
buffer queues is associated, for each of the time frames with a unique combination of one
of the incoming channels and one of the outgoing channels, wherein each of the buffer
5 queues is further comprised of an alignment subsystem comprised of a plurality of time
frame queues;

logically mapping, for each of the (UTR-j) time frames selected incoming
channels (j) to selected buffer queues;

10 logically mapping, for each of the CTR time frames, selected ones of the
plurality of buffer queues to selected outgoing channels (l),

determining when the respective time frame queue is empty;

determining when the respective time frame queue is not empty;

15 storing the data packets that arrive via the incoming channel (j) in the
respective time frame queue of the alignment subsystem responsive to the logically
mapping, and

coupling selected ones of the time frame queues to respective ones of the
outgoing channels (l), for transfer of the respective stored data packets during the
respective associated CTR time frames.

008220-103800

71. The method as in claim 70,
transferring all of the data packets associated with a respective first time frame into an empty first time frame queue from incoming channel (*j*), during the respective selected first time frame of the time frames (TFs) as was defined by UTR-*j*,
wherein the respective time frame queue is designated as full, responsive to the logically mapping;

72. The method as in claim 71, further comprising maintaining the first time frame queue and the second time frame queue as mutually exclusive at all times.

74. The method as in claim 69, further comprising:
providing an optical link with a plurality of optical channels as the
communications link; and
adding a delay element to a selected one of the input ports.

5

10

15

20

- 119 -

where each of the communications links has a plurality of channels, each associated with a respective wavelength.

80. The method as in claim 79, further comprising:

5 coupling a first predefined subset of the channels for each respective one of the communications links from the respective communication link to a second defined one of the communications links.

81. The method as in claim 80, further comprising:

10 demultiplexing the first predefined subset into a predefined respective second predefined subset of the respective channels.

82. An input interface method comprising:

15 mapping an asynchronous stream of data packets, each comprising a header portion and a payload portion, from at least one source to at least one destination;

providing a Common Time Reference (CTR), divided into a plurality of contiguous periodic super cycles each comprised of at least one contiguous time cycles each comprised of at least one contiguous time frame (TF);

20 providing at least one synchronous virtual pipe (SVP) having a subset of predefined ones of the time frames uniquely associated therewith;

analyzing the header portions of the asynchronous data packets;

providing a link coupled to the destination;

10 identifying respective ones of the data packets as CBR data packet, VBR
data packet, and BE data packet;

15 forwarding a respective one of the data packets from the respective one of
queues that is associated with the respective time frame responsive to the CTR.

prioritizing the output from the respective one of the queues to provide

1756 Draft Appl.v1 (3/28/00)

respective one of the queue's VBR queue, and then from the respective one of the queue's
BE queue.

84. The method as in claim 83, further comprising:

5 determining which certain ones of the data packets from the CBR, VBR,
and BE queues for the respective one of the time frames are not output during a
respective associated one of the time frames; and

detecting and rescheduling the certain ones of the data packets.

10 85. The method as in claim 82,
wherein each SVP is associated with at least one Pipe ID (PID).

86. The method as in claim 85,

15 wherein the PID is one of the following: explicitly contained in a field of
the data packet header portion, implicitly given by an Internet protocol (IP) address,
Internet protocol group multicast address, a combination of values in the IP address and
transport control protocol (TCP), a user datagram protocol (UDP) header, an MPLS
label, an asynchronous transfer mode (ATM) virtual circuit identifier (VCI), and an ATM
virtual path identifier (VPI), a combination of VCI and VPI.

20

5

10

15

20

[illegible]